Applying Genetic Programming to Economic Mechanism Design: Evolving a Pricing Rule for a Continuous Double Auction

Steve Phelps
Peter McBurney
Dept of Computer Science
University of Liverpool
Chadwick Building
Liverpool L69 7ZF, UK.
sphelps,peter@csc.liv.ac.uk

Simon Parsons
Dept of Computer and Information Science
Brooklyn College
2900 Bedford Avenue
Brooklyn, NY 11210, USA.
parsons@sci.brooklyn.cuny.edu

Elizabeth Sklar
Dept of Computer Science
Columbia University
1214 Amsterdam Avenue
New York, NY 10027, USA.
sklar@cs.columbia.edu

Categories and Subject Descriptors

1.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence: Multi-agent systems

General Terms

Design, Economics

Keywords

cognitive game theory, continuous double auction, genetic programming, pareto optimisation, mechanism design, multi-agent systems (MAS), reinforcement learning, trading strategies

1. INTRODUCTION

The auction mechanism design problem has attracted much interest in recent years, and economists have had considerable success in applying techniques from game theory to the design of auction-based markets for deregulated commodity markets (e.g., California's deregulated electricity market) and the sale of government assets (e.g., auctions of electromagnetic spectrum for mobile phones). Alvin Roth has suggested that this is akin to an engineering process in which economists design the rules of a market mechanism in order to meet particular socio-economic requirements (e.g., maximising the efficiency of allocating commodities in a market).

The engineering of auction mechanisms is of particular importance to agent-based electronic commerce and multi-agent systems in general. E-commerce has enabled consumers to act as price-makers instead of just price-takers in large auction-based markets and has stimulated the use of personalised bidding agents to empower those consumers even more. In addition, auction mechanisms are seen as a promising means of solving many distributed resource-allocation problems in multi-agent systems and grid technology.

One approach to mechanism design is to use techniques from machine learning to explore the space of possible ways in which agents might act in particular markets. For example, reinforcement learning has been used to explore bidding patterns in auctions and to establish the ways in which price-setting behavior can affect consumer markets. Another approach is to use techniques from evolutionary computing, e.g., co-evolutionary machine-learning. Our earlier work has explored the use of co-evolutionary GP to determine auction mechanism rules automatically.

In that work, mechanism rules and bidding strategies were co-evolved in ways that sought to maximise, on the one hand overall market efficiency, and on the other hand the profits of individual agents. This approach lends itself to studying the dynamics of the evolution of negotiation mechanisms in a setting where mechanisms change incrementally in a real environment, but it is problematic when we wish to derive *optimal* mechanism designs

In our later work we view mechanism design as a multi-objective optimisation problem. The key issue that we address is determining the fitness of individual points in the mechanism design-space. This is non-trivial in the general case, since assessing the fitness of an individual mechanism involves reasoning about how agents might actually behave under the proposed negotiation rules. In the following section we describe our approach to predicting agents behaviour for arbitrary mechanisms, and in the final section we present a summary of our results where we apply this method to mapping the fitness landscape for a k-CDA.

2. EQUILIBRIA FOR N-PLAYER GAMES

When evaluating a mechanism design, the designer must take into account the set of trading strategies that are likely to be played by agents trading in the mechanism under consideration. Deriving the set of the strategies likely to be played for a particular market game, that is "solving" the game, is a non-trivial problem in the general case. This is because there is often no clear *dominant* strategy which

AAMAS'03, July 14–18, 2003, Melbourne, Australia.

constitutes best play; rather the best strategy to play depends entirely on the strategies played by other agents. Nash defined a solution concept in which the strategy adopted by any given agent is a best-response to the best-response strategies adopted by all other agents, and proved that all n-player, non-zero-sum games admitted solutions so defined.

Nash's solution concept is widely adopted in theoretical economics. Thus when evaluating an economic mechanism, the designer computes the Nash equilibria of strategies for the given mechanism; and this forms the basis of predictions about how people will actually behave under the rules of this mechanism. The designer can then analyse market outcomes in equilibria and quantitively assess, for example, the likely affect on overall market-efficiency that a given change in the mechanism rules will yield. Thus the role of the designer is to ensure that the Nash equilibria correspond to situations in which high market efficiency is obtained.

We can view mechanism design as a *multi-objective optimization* problem. We consider as a separate dimension each problem variable we are interested in maximising (for example, market efficiency, seller revenue and so on), and the difficulty lies in simultaneously maximising as many dimensions as possible. The designer's task is to choose mechanism rules which pareto-optimise different market variables when traders play Nash-equilibrium strategies. However, there are a number of problems beginning with computing the Nash equilibria:

1. Agents with limited computational power (i.e., "bounded rationality" constraints) may be unable to compute their Nash-equilibrium strategy;

2. Even with vast amounts of computational and analytic power, many games defy solution; e.g., in the case of the k-double-auction, analytical techniques have yet to yield a solution;

3. Empirical evidence shows that human agents often fail to coordinate on Nash-equilibria for very simple games whose solution is easily derivable under bounded-rationality assumptions; and

4. Often a given game will yield a multitude of Nash solutions and there is little guidance for practitioners on choosing plausible subsets thereof as predicted outcomes.

These difficulties with the standard theory of games have led to the development of a field known as *cognitive game theory*, in which models of learning play a central role in explaining and predicting strategic behaviour. Erev and Roth show how simulations of agents equipped with a simple reinforcement learning algorithm can explain and predict the experimental data observed when human agents play a diverse range of trading games. Such *multi-agent reinforcement learning* models form the basis of our solution concept for optimising mechanism designs. Rather than computing the theoretical equilibria for a given point in the mechanism search space, we run a number of multi-agent simulations using agents equipped with a learning algorithm that determines their bidding strategies. The stationary points in these simulations – the states where the learning algorithms of all agents have converged – correspond to the equilibria of classical game theory, and the market outcomes in these stable states can be viewed as predictions.

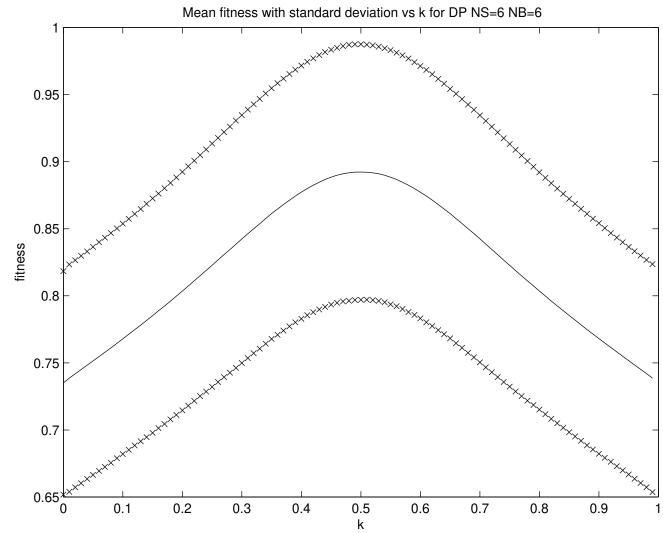

Figure 1: Fitness plotted against k for a discriminatory-price k-CDA with 6 buyers and 6 sellers

Note that we are not attempting to find theoretically optimal strategies for our agents[1]. Rather, we are attempting to predict how boundedly-rational agents, who have no prior knowledge of an equilibrium solution nor the means to calculate one, might actually play against the mechanism we are (automatically) designing. For this reason, we chose to use the Roth-Erev algorithm , since it forms the basis of a *cognitive model* of how people actually behave in strategic environments. In particular it models two important principles of learning psychology:

- *Thorndike's law of effect* — choices that have led to good outcomes in the past are more likely to be repeated in the future; and

- *The power law of practice* — learning curves tend to be steep initially, and then flatter.

3. RESULTS

Figure 1 shows the mean fitness of a set of k-CDA mechanisms for 100 values of k in the interval [0, 1] at intervals of 0.01. Each sample consisted of 100,000 runs of the auction simulation with different seeds for the random number generator. In each simulation we pit 6 buyers against 6 sellers over a period of 1,000 rounds of trading. Fitness is defined as a linear sum of buyer market-power, seller market-power and overall market efficiency, normalised to lie in the range [0, 1] where higher values indicate better mechanisms. The software used to run this experiment is available for download at:

`http://www.csc.liv.ac.uk/~sphelps/jasa`.

Also available at the same URL is the full version of this paper, in which we demonstrate the evolution of a k=0.5 CDA using genetic programming.

[1]In other words Nash equilibrium strategies